
Theme: Harvard-Riken Joint Symposium :

Application of GPU Computation to Brain Science, Quantum Science, Astronomy, Fluid Dynamics and other sciences

Date: Aug. 28th (Fri) 9:30-18:10; Poster Session 13:00-14:00; Reception 18:30-20:00

Aug. 29th (Sat) 9:30-

Place: Ohkouchi Hall, RIKEN (The Institute of Physical and Chemical Research)

Organizers: Toshiaki Iitaka, Toshikazu Ebisuzaki, Ryutaro Himeno, Makoto Taiji

Michael Stopa, Alan Aspuru-Guzik

Franco Nori, Koji Ishibashi

Contact: Toshiaki Iitaka

E-mail: hariken09@riken.jp

[Program] [Satellite Seminars][Photo][Abstract]

http://www.iitaka.org/hariken09.html#program

http://www.iitaka.org/hariken09.html#seminar

http://picasaweb.google.co.jp/HaRiken.Collaboration/20090828HaRiken09#slideshow/

http://www.iitaka.org/hariken09_abstract.pdf

Links: GPGPU http://www.iitaka.org/gpgpu.html

Computation with GPU http://en.wikipedia.org/wiki/Gpgpu

sciGPU.org http://www.scigpu.org/

RICC http://w3cic.riken.go.jp/ricc/index_e.html

HaRiken08 http://www.iitaka.org/hariken08.html

HaRiken08 http://sites.google.com/a/aspuru.com/hariken-08/Photo

Access to RIKEN:

http://www.brain.riken.jp/en/visiting/index.html#international

http://www.riken.go.jp/engn/r-world/riken/campus/wako/index.html

Hotel

http://www.toyoko-inn.com/e_hotel/00090/index.html

Easiest route is to use direct Bus from Narita to Wako.

Hotel is 5min from the Bus stop.

==========

# 理研ハーバード大学連携シンポジウム：
# 脳科学、量子科学、天文学、流体力学等への GPU 計算の応用

このページは随時更新しています。[再読み込み]をして最新版をご覧くださ い。
プログラムの最新情報はこのページに掲載されます。

http://www.iitaka.org/hariken09.html

開催日時　　：　　2009 年 8 月 28 日(金) 9:30-18:10;　ポスターセッション　13:00-14:00　**レセプション：　18:30-20:00**

29 日（土）　9:30-

場所　　：　　理化学研究所　大河内記念ホー ル

主催　　：　　理化学研究所　基幹研究所

実行委員：　　飯高敏晃、戎崎俊一、姫野龍太郎、泰地真弘人

Michael Stopa, Alan Asupru-Guzik

Franco Nori, 石橋幸治

問合先　：　　戎崎計算宇宙物理研究室　飯高敏晃

E-mail:　hariken09@riken.jp

URL: http://www.iitaka.org/

概要：

GPU 計算の脳科学、量子科学、天文学、流体力学等への応用を議論します。シンポジウムは公開形式とし，幅広い立場からの活発な議論を期待します。

参考リンク：　GPU 計算とは。 http://ja.wikipedia.org/wiki/GPGPU

sciGPU.org　http://www.scigpu.org/

RICC　http://www.riken.go.jp/r-world/info/info/2009/090807/index.html

HaRiken08 http://www.iitaka.org/hariken08.html

HaRiken08 http://sites.google.com/a/aspuru.com/hariken-08/Photo

東京大学物性研 究所 短期研究会 計算物理学 (GPGPU 特別セッション) 日時: 2009 年 12 月 10 日(木) 10:00 〜12 月 11 日(金) 16:20(アブストラ クト)

http://www.issp.u-tokyo.ac.jp/public/keisan09/

http://www.issp.u-tokyo.ac.jp/public/keisan09/abstract.pdf

==========

# Program

## Aug. 28th (Fri)  9:30-18:10

- Opening:  Toshiaki Iitaka (RIKEN)
  - 9:30- 9:40 Kohei Tamao (RIKEN)
  - 9:40- 9:50 Michael Stopa PostCard (Harvard & RIKEN)
  - 9:50-10:00 Koji Kaya (RIKEN)
- Brain Science: Hideyuki Cateau (RIKEN) (RIKEN)
  - 10:10-10:50 Visualizing Columnar Functional Architectures in Humans Using High-Field High-Resolution Functional Magnetic Resonance Imaging
    Kang Cheng (RIKEN)
  - 10:50-11:30 Scalable and Interactive Segmentation and Visualization of Neural Processes in EM Datasets
    Won-Ki Jeong (Harvard) (Connectome)
- Nano Devices:  Keiji Ono (RIKEN)
  - 11:30-12:00
    Koji Ishibashi (RIKEN)
  - 12:00-12:30 GPU for designing Semiconductor devices
    Michael Stopa (Harvard & RIKEN)

## Lunch 12:30-14:00

- Quantum Science:  Michael Stopa (Harvard & RIKEN)
  - 14:00-14:30 GPU-accelerated Computing for Earth and Planetary High Pressure Science
    Toshiaki Iitaka (RIKEN) (GPGPU) (Paper)
  - 14:30-15:10 General-Purpose GPU computing for Quantum Chemistry
    Alan Aspuru-Guzik (Harvard)
  - 15:10-15:40  Brief pedagogical overview on quantum computing using superconducting qubits
    Franco Nori (RIKEN)

Coffee 15:40-16:00

- Astronomy & Fluid Dynamics: Makoto Taiji (RIKEN)
  - 16:00-16:40 High Performance Computing as Lynchpin in Next-generation Radio Telescopes
    Lincoln Greenhill (Harvard) (MWA)
  - 16:40-17:10 Particle simulations with GPU
    Naohito Nakasato (Aizu Univ. & RIKEN)
  - 17:10-17:40 The use of GPU for practical astrophysical N-body simulations
    Keigo Nitadori (RIKEN)
  - 17:40-18:10 Performance evaluation of himenoBMT on Riken Integrated Cluster of Cluster Including GPGPU and GUI Code Developing Environment for GPGPU
    Ryutaro Himeno (RIKEN)
- 18:30-20:00 Reception at Hirosawa Club

## Poster Session 13:00-14:00

(you can post 12 pieces of A4 paper or 1 piece of A0 portrait paper)

- Accelerating Orbital-Free First Principles Calculation with Graphics Processing Unit
  Masaru Aoki (Shizuoka Sangyo Univ)
- MEM charge density analysis by using GPU
  Hiroshi Tanaka (Shimane Univ.)
- GPU Acceleration of the Diffusion Monte Carlo Method
  Teppei Ono (Univ. Tokyo)
- Acceleration of Fragment MO method using CUDA on GPU
  Yuki Furukawa (X-Ability Co., Ltd)
- Computation of atmospheric delay corrections on GPUs
  Thomas Hobiger (NICT)
- GUI Code Developing Environment for GPGPU on RICC : RIKEN Integrated Cluster of Clusters
  Shigeho Noda
-
  Koji Yasuda (Nagoya Univ.)

## Aug. 29th (Sat)　9:30-

- GPU computing and beyond:　Lincoln Greenhill (Harvard)
  - 9:30-10:00 GPU Based Acceleration of First Principles Calculation
    Hidekazu Tomono　(Meiji Univ.)
  - 10:00-10:30 A real-time GPS software receiver running on the GPU
    Thomas, Hobiger　(NICT)
  - 10:30-11:00 Multi-precision linear equation solver for lattice QCD
    Michael Clark　(Harvard & Boston)
- Quantum Chemistry:　Alan Aspru-Guzik (Harvard)
  - 11:10-11:40 A mixed-precision matrix multiplication library for GPUs and its application to Quantum Chemistry calculations
    Mark Watson　(Harvard)
  - 11:40-12:10 GPU acceleration of ab-initio quantum Monte Carlo simulations and its application to molecular crystals
    Kenta Hongo　(Harvard)
  - 12:10-12:40 Computational Aspects of ab-initio QMC calculation
    Ryo Maezono　(JAIST)

# Participants without presentation

Tomoharu, Terashima　(JAIST)
Masanori Tachikawa　(Yokohama City Univ.)
Isao Kawata　(Canon Inc.)
Shinichiro Nakamura　(Mitsubishi Chem. Inc.)
Tadahiro, Gotoh　(NICT)
Hiroshi Murakami　(Tokyo Metropolitan University)
Takeshi Yamamoto　(Kyoto Univ.)
Makoto, Matsuyama　(Best Systems Inc.)
鷲田　充宏　(Research Center of Computational Mechanics, Inc)
荒川　貴道　(Research Center of Computational Mechanics, Inc)
三又　秀行　(Research Center of Computational Mechanics, Inc)
古賀　良太　(X-Ability Co., Ltd)
古川　祐貴　(X-Ability Co., Ltd)
西村　涼平　(X-Ability Co., Ltd)

Hideo, Tsuru (NITTOBO ACOUSTIC ENGINEERING CO.LTD )
Jae Duck, Lim (Hynix Semiconductor Japan)
大村一太 (理研高速分子シミュレーションチーム)
杉原崇憲 (理研次世代スパコン)
田代英俊 (日本科学技術振興財団)
長坂由子 (日立製作所中央研究所)
池田弘行 (東芝　生産技術センター)
三浦均 （武蔵野美術大学　映像学科）

参加（発表）申し込み：
下記の様式をメールにコピーして記入のうえ 8 月 10 日から 8 月 21 日午前 までに hariken09@riken.jp にお申し込み下さい。
申込多数の場合はそれ以前に締め切ります。
==========
==========
参加申込：
氏名：
所属：
住所：
電話番号：
電子メールアドレス：
参加人数：　____人
レセプション参加人数（有料 3,000 円／人）　X ____人
発表申込：
ポスター発表と少数の口頭発表を募集します。希望者は以下の情報を英語で加えてください。
type of presentation: poster/oral
title:
name (affiliation) of the authors:
abstract: less than 1/2 page
e-mail:
==========
==========
個人情報に関する取扱について：
ご記入いただいた個人情報は，参加者の集計および主催者からの連絡に利用させていただきます

Please cut&paste the "Application Form" into text email, fill it and send it to hariken09@riken.jp.
Submission Period: 10th-21th Morning, Aug. (but it will be closed earlier when many applications are submitted.)

==========

Application Form:

name:
affiliation:
address:
telephone:
e-mail:
number of participants: ___ person
reception (3,000 yen/person) x ____ person

Presentation ( optional ):

type of presentation: poster/oral
title:
name (affiliation) of the authors:
abstract: less than 1/2 page
e-mail:

==========

# Satellite Seminars

- Subject : Accelerating Quantum Chemistry calculations using graphical processing units (GPUs)

Speaker : Dr. Mark Watson (Harvard University, USA)
Language : English
Date : Monday, Sept. 7, 2009 11 : 00 ～ 12 : 30
Location : Seminar Room (224,226), 2F Main Building

Contact : Next-generation Molecular Theory Unit

- August 27 (Thu) 10:30 am- 11:30 am
  4th Floor Seminar Room (#435) Main Research Building
  SPEAKER: Dr. Alán Aspuru-Guzik (Department of Chemistry and Chemical Biology, Harvard University)
  TITLE: Quantum Computation for Quantum Chemistry
  contact: Digital Materials Team

## Aug. 28th(Fri) 9:30-12:30

## Opening: Kohei Tamao (RIKEN)

Welcome Address 20090828

Good morning ladies and gentlemen.
I am Kohei Tamao, the director of RIKEN Advanced Science Institute ASI. Welcome to RIKEN Wako campus in the worst season, hot summer. I clearly remember that when I visited Harvard University in March this year, we had an unseasonal snowstorm; it was frosty cold outside. In contrast, here in RIKEN, it is hot outside today, so it may be a good idea for you to stay inside this comfortable meeting hall to enjoy science.

Now, on behalf of the ASI, it is a pleasure for me to give a welcome address to the Harvard-Riken Joint Symposium, so-called HaRiken09, entitled "Application of GPU Computation to Brain Science, Quantum Science, Astronomy, Fluid Dynamics and other sciences", co-organized by Drs. Iitaka, Ebisuzaki, Himeno, and Taichi in Riken side and Dr. Mike Stopa and Dr. Alan Aspuru-Guzik in Harvard side.
So, I would like to express my sincere thanks first to the organizers for their efforts and also all the participants who have come from not only inside Japan but also from overseas.
Especially, it is a pleasure to have some leading researchers and experts in this field with us today, who include
Dr. Michael Stopa,
Dr. Aspuru-Guzik,
Dr. Lincoln Greenhill, and
Dr. Won-Ki Jeong.
Thank you for traveling a long way to visit RIKEN.

Since our institute RIKEN Advanced Science Institute ASI has a pleasure to support financially this symposium, I would like to make a brief introduction to our institute, ASI on this occasion.
The ASI, as RIKEN's core organ, was established last year by merger of two prior institutions, the Discovery Research Institute and the Frontier Research System, which had different important missions, respectively, the germination of new seeds of science by curiosity-driven basic research and the incubation of these new ideas to new fields of

science and technology by integrated, strategic research.
Thus, the philosophy and mission of the ASI is summed up in one phrase; Vital Integration for Emergence and Innovation.
We have about 700 full-time researchers; this is about 20% of the total number of researchers in RIKEN, covering all fields of natural sciences, physics, chemistry, engineering, biology and life-sciences.
The ASI budget including tenured employees is about 9 billion yen, about 90 million US dollars, corresponding to about 10% of RIKEN's total budget.
The mission of our institute is defined to create new research fields by interdisciplinary interactions among researchers in and outside of RIKEN. One of the most important tasks of ASI is the promotion of international collaborations of scientific research. As part of our globalization strategy, we support financially some important international symposiums like this.

Let's back to this symposium theme. The ASI has been promoting research collaboration with Harvard University with Dr. Yasunori Yamazaki as the leader. One of the research fields is Application of GPU Computation, the theme of this symposium.
I understand that this collaboration was initiated in early 1990s when Mike Stopa san visited RIKEN and stayed for several years to do collaborative research with Iitaka, and Ishibashi in the RIKEN Frontier Research System, one of the prior organizations of the ASI. In 2006, Mike and Iitaka proposed a RIKEN-Harvard collaborative research project to the President funds. The idea was pursued strongly together with Dr. Kaya, former director of Discovery Research Institute, the other partner of the ASI, and Franco Nori and others. In 2008, the first RIKEN-Harvard international symposium on the GPU computation, HaRiken08, was held at Harvard by Dr. Aspuru-Guzik, Dr. Stopa and Dr. Iitaka, to which the RIKEN ASI sent a delegation of several leading scientists. The symposium today HaRIKEN09 is thus the second Harvard-RIKEN symposium on this topics.

Since I am a chemist, synthetic organic chemist, I am not familiar to the GPU computation. Iitaka san kindly lectured me how the GPU computation is an exciting science. Now, I find that the processing speed of GPU is 100 times faster than that of the current CPU for appropriate subjects and algorisms. I have been so much excited by his lecture. If their approach is accomplished to mount the GPU to personal computers, the calculation speed of ordinary PCs can become comparable to that of current super computers. We can calculate a variety of large-scale and highly-detailed simulations

and analysis of experiments with desktop PCs. It is amazing and it makes dramatic advances in and great impact on science and technology. I hope that this can be realized in the near future.

I have just heard a hot news that a group of RIKEN researchers lead by Dr. Nitadori has achieved 42 trillion floating point calculations per second by using GPU cluster consisting of 256 GPU and they are selected as one of the finalists of Gordon Bell Prize 2009, known as the Nobel Prize in the computational science field. I am proud of such able young scientist at RIKEN and I really hope that they can win the gold medal.

Finally, I really hope that the present symposium aiming for the application of GPU computation to brain science, quantum science, astronomy, and fluid dynamics is fruitful and will contribute to innovative achievements of the computational science and technology. Please enjoy the symposium.
Thank you.

# Oral

"Visualizing Columnar Functional Architectures in Humans Using High-Field High-Resolution Functional Magnetic Resonance Imaging"

Kang Cheng (RIKEN)

Since its inception early in 1990s, functional magnetic resonance imaging (fMRI) has become one of the most dominant neuroimaging tools for studying human brain fucntions. High-resolution fMRI, with its improved signal-to-noise ratio and spatial specificity, has strengthened the capability of fMRI and allowed mapping of functional architectures in human brains. In this talk, I will first explain the principle of the blood oxygenation level-dependent (BOLD) effect, upon which most of fMRI experiments are conducted, factors restricting the spatial specificity of the BOLD signal, and measures dealing with these factors. I will then briefly introduce several practical techniques developed using high-field MRI systems for revealing functional architectures in human visual cortices.

# “Scalable and Interactive Segmentation and Visualization of Neural Processes in EM Datasets”

Won-Ki Jeong (Harvard,Connectome)

Recent advances in scanning technology provide high resolution EM (Electron Microscopy) datasets that allow neuroscientists to reconstruct complex neural connections in a nervous system. However, due to the enormous size and complexity of the resulting data, segmentation and visualization of neural processes in EM data is usually a difficult and very time-consuming task. In this talk, I will present NeuroTrace, a novel EM volume segmentation and visualization system that consists of two parts: a semi-automatic multiphase level set segmentation with 3D tracking for reconstruction of neural processes, and a specialized volume rendering approach for visualization of EM volumes. It employs view-dependent on-demand filtering and evaluation of a local histogram edge metric, as well as on-the-fly interpolation and ray -casting of implicit surfaces for segmented neural structures. Both methods are implemented on the GPU for interactive performance. NeuroTrace is designed to be scalable to large datasets and data-parallel hardware architectures. A comparison of NeuroTrace with a commonly used manual EM segmentation tool shows that our interactive workflow is faster and easier to use for the reconstruction of complex neural processes.

“Nanotubs and nanowires for building blocks of nanodevices”

Koji Ishinashi (RIKEN)

There has been increasing interest in the nanotubes and nanowires that are formed in a self-assembled manner with a diameter that is difficult to realize with a conventional lithography technique.

We show electrical transport of single-wall carbon nanotubes at low temperatures, and show that they exhibit artificial atom behaviors. This include the simple shell structures, and Zeeman spitting of one particle states and singlet-triplet difference when two electrons in a shell. Another interesting fact is that the energy scale associated with the artificial atom fall into the THz range, for which we could observe Thz photon assisted tunneling. We believe these behaviors are coming from the one-dimensional nature of the carbon nanotubes, and we show our recent effot to observe similar effect with semiconductor nanowires that are grown by, so called, VLS method.

“GPU for Modeling Semiconductor Devices”

Michael Stopa (Harvard ,RIKEN)

I will describe self-consistent electronic structure calculations for semiconductor devices. In particular I will concentrate on the SETE code for GaAs-AlGaAs two-dimensional electron gas heterostructure based quantum devices; specifically quantum wires and quantum dots. These structures are used as possible candidates for future electronics and quantum electronics. The algorithm for calculating the electronic properties of such devices, embodied in SETE, is based predominantly on sparse matrix linear algebra. For quantum dots, three portions of the code in particular represent bottlenecks for processing: the Poisson equation, the reduction of the 3D calculation of the wave functions to an effective 2D calculation, and the eigenvalue problem. I will describe these portions of the code and the prospect of designing CUDA kernels to parallelize these functions on the GPU.

## Aug. 28th(Fri) 14:00-18:10
## Oral

"GPU-accelerated Computing for Earth and Planetary High Pressure Science"

Toshiaki Iitaka (RIKEN)

I will introduce what kind of problems in Earth and Planetary High Pressure Science can be accelerated by GPU computing. In particular, I will discuss the quantum molecular dynamics in tight binding approximation and the spin-1/2 (qubit) dynamics.

[1] http://uk.arxiv.org/abs/0910.4497
[2] http://www.iitaka.org/gpgpu.html

“General-Purpose GPU computing for Quantum Chemistry”

Alan Aspuru-Guzik (Harvard)

In this talk, I overview my group and collaborator's efforts towards the acceleration of quantum chemistry calculations using general-purpose graphics processing units for the acceleration of quantum chemistry calculations. In particular, I will discuss our efforts in the acceleration of second-order Moller-Plesset perturbation theory (MP2) using the resolution of the identity (RI) approximation, as well as real-space quantum Monte Carlo (QMC). I will also describe the efforts of Q-Chem, Inc. towards accelerating localized-basis set density functional theory calcualtions. The development of general-purpose tools to aid with the porting of quantum chemistry packages to GPGPU will be discussed.

“Brief pedagogical overview on quantum computing using superconducting qubits”

Franco Nori (RIKEN)

Superconducting (SC) circuits can behave like atoms making transitions between a few energy levels. Such circuits can test quantum mechanics at macroscopic scales and be used to conduct atomic-physics experiments on a silicon chip. This presentation overviews a few of our theoretical studies on SC circuits and quantum information processing (QIP) including: SC qubits for photon generation and for lasing; 2-1 photon coexistence; cooling qubits and their environment; using SC qubits to probe nearby defects; hybrid circuits involving both charge and flux qubits; quantum tomography in SC circuits; preparation of macroscopic quantum superposition states of a cavity field via coupling to a SC qubit; generation of nonclassical photon states using a SC qubit in a microcavity; and controllable couplings among qubits. These controllable couplings between qubits can be achieved either directly or indirectly [for instance, with or without coupler circuits; as well as with or without "data-buses" like resonators (e.g., electromagnetic fields in cavities, LC circuits, or transmission line resonators)].

For brief overviews, see: J.Q. You and, F. Nori, Physics Today 58, No. 11, 42 (2005).

F. Nori, "Atomic physics with a circuit", Nature Physics 4, 589 (2008).

“High Performance Computing as Lynchpin in Next-generation Radio Telescopes”

Lincoln Greenhill (Harvard,MWA)

Frontier astronomical facilities in the next decade will pose serious data-intensive science challenges. In a new paradigm for radio astronomy, High Performance Computing will be front and center as a critical element of the interferometric arrays that will make tomographic maps of the early universe. Power-efficient green implementation will be an additional requirement for some projects.

Computation using Blue Gene-class clusters, mass deployment of Field Programmable Gate Arrays, and Graphics Processing Units figure prominently in array plans and prototypes, though scaling to the very largest radio array now in development, planned for 2020, is an unsolved problem. Lessons learned and preliminary images from one of the first path-finding instruments will be presented.

“ Particle simulations with GPU”

Naohito Nakasato  (Aizu Univ. & RIKEN)

We will present our implementation of particle simulations (N-body) accelerated with GPU. Especially, we present a technique to implement O(N log N) algorithm, which is much more efficient than a brute-force O(N2) algorithm, for computation of force between particles.Our technique apply to not only simple gravity interaction but also complex particle interactions such as Smoothed Particle Hydrodynamics method and general neighbor force.

“The use of GPU for practical astrophysical N-body simulations”

Keigo Nitadori (RIKEN)

Gravitational many body problem has been an important target of General Purpose GPU (GPGPU) from its early days, mainly because of its simpleness, rich parallel nature or small requirement on memory access. After the release of unified shader architecture GPU followed by CUDA, a number of computer scientists or astrophysicist have demonstrated GPU $N$-body implementations at a level of a few hundreds Gflops on single GPU card. However, these pioneering works are based on brute $O(N^2)$ algorithm and single precision arithmetics, and cannot be directly applied for practical astrophysical simulations.

For the practical calculations, more sophisticated algorithm or higher precision is needed. For example, for the simulations of collisional systems (planetesimal systems, star clusters or galactic nuclei), temporal hierarchical algorithm called 'individual timestep scheme' is widely used with direct summation method for the force calculation. Double precision accuracy is required for (not all but) some parts of calculations.

For the simulations of collisionless systems (galaxies, clusters of galaxies or cosmological structure formations), single precision supplies an enough accuracy, but $O(NlogN)$ or $O(N)$ spatial hierarchical approximation algorithms are used and some efficient uses of GPUs for such algorithms are to be considered (or the effective performance drops down to the same level of well optimized CPU code exploiting multi-core and SIMD instructions).

We discuss on some new approaches on the use of GPU for practical the Simulations of collisional/collisionles systems. For the collisional simulations, we introduce a 'Pseudo double precision algorithm' which emulates a double precision arithmetic with a combination of some single precision arithmetics. Note that the fraction of the precision sensitive part in the entire calculation is small and for the most of them can be executed in single precision arithmetics. Relative overhead of introducing the pseudo double precision is less than $¥sim 30 ¥%$.

For the collisionless simulations, we present a modified 'a modified Barnes-Hut tree algorithm for GRAPE and vector processors' algorithm for GPUs. The basic idea of the new method is to map a new parallelism, which has been executed serially in the previous mtthod to the multiprocessors of the GPU. We achieved a 3x speed-up with single G92 GPU compared to a well optimized CPU implementation for SSE and quad-core.

“Performance evaluation of himenoBMT on Riken Integrated Cluster of Cluster Including GPGPU and GUI Code Developing Environment for GPGPU”

Ryutaro Himeno (RIKEN), Akira Naruse(Fujitsu Laboratory ),
Takayuki Shigetani(RIKEN), Shigeho Noda(RIKEN), Hideaki Komatsu (IBM)

We installed a new system called “RIKEN Integrated Cluster of Cluster” with 100 GPGPU boards(NVIDIA Tesla) in July, 2009 and have just started operation. Before starting service, we measured himenoBMT as well as LINPACK benchmark test. The himenoBMT is a core of 3-d pressure Poisson solver in an incompressible flow solver using FDM method on curvilinear coordinate system wrote by one of the authors. Dr. Naruse, Fujitsu Laboratory has modified the himenoBMT code on GPGPU and improved the performance. We have checked and measured the performance of the himenoBMT code on the cluster with Tesla. We will show its results.

On the other hand, code development highly adjusted to GPU is generally very difficult for users. ACCC was installed the GUI development environment develop by IBM Japan and both IBM Japan and RIKEN started joint research of its new applications. We will show it in detail in the workshop.

# Poster

"Accelerating Orbital-Free First Principles Calculation with
Graphics Processing Unit"

Masaru Aoki (Shizuoka Sangyo Univ.) ,Hidekazu Tomono(Meiji Univ.),
Toshiaki Iitaka(RIKEN), Kazuo Tsumuraya(Meiji Univ.)

Computational material design requires efficient algorithms and high-speed computers for calculating and predicting material properties. The orbital-free first principles calculation (OF-FPC) method, which is a tool for calculating and designing the properties, is an O(N) method and is suitable for large-scaled systems. In this study, we implement CUFFT routine, which is a FFT library of CUDA for GPGPU (General-Purpose Graphics Processing Unit), into our in-house OF-FPC code. The implementation of the CUFFT reduces the total computation time to half of that of the CPU.

“MEM charge density analysis by using GPU”

Hiroshi Tanaka (Shimane Univ.), Tomonori Hattori (Materials Science, Shimane Univ.)
Toshiaki Iitaka (RIKEN),Masaki Takata (RIKEN/SPring-8)

We have been developed a program package, which evaluates electronic charge density in detail for huge crystalline systems, such as polymers and proteins, from the x-ray diffraction data by using the maximum entropy method (MEM) [1]. In this study, we tried to utilize a GPU, because this type of analysis needs a lot of CPU resources.

A preliminary result shows that the elapsed time for the MEM analysis by using GPU is 4 to 5 times as fast as that by only CPU. It is expected that the speedup ratio finally achieves to more than 10 after tuning the program.

[1] H. Tanaka et al. J. Appl. Crystallogr. Vol. 35. 282 (2002).

“GPU Acceleration of the Diffusion Monte Carlo Method”

Teppei Ono (Univ. Tokyo), Shinji Tsuneyuki(Univ. Tokyo)

First-principles electronic structure calculation requires a lot of computing resources to calculate large systems such as proteins, molecular crystals, and the concern with the GPGPU has been growing.

Focusing on super parallel processing of GPU, we are developing codes of variational Monte Carlo (VMC) and diffusion Monte Carlo (DMC) methods on GPU. These methods are relatively simple to parallelize and are regarded as promising methods of investigating molecules and solids with highest accuracy.

We have succeeded in calculating total energy of atoms by VMC and DMC methods on GPU. Our codes provide results of higher precision than single precision by partially using double precision computing. The VMC calculation on GPU is about 5x faster than those on CPU in case of Ar atom. In general, efficiency of parallelization improves with system size, so that further speed-up is expected with larger systems.

In this presentation, we discuss efficient algorithms of VMC and DMC methods, and show details of the method of high precision calculation.

# “ Acceleration of Fragment MO method using CUDA on GPU”

Yuki Furukawa (X-Ability Co., Ltd) ,Ryohei Nishimura(X-Ability Co., Ltd),
Ryota Koga(X-Ability Co., Ltd)

Recently, the methods of computational chemistry become more applicable to macromolecules such as proteins with remarkable improvement of computational resources. 2. General Purpose GPU (GPGPU) is successfully applied for various problems in the field of high-performance computation using NVIDIA CUDA Environment. 3. GPGPU is already applied for some ab initio quantum chemistry calculations, but FMO method, which has high parallelization efficiency, remains untouched in GPGPU field.

We developed FMO-MP2 algorithm suitable for GPU with CUDA environment. The performance is compared with GAMESS 2009 Jan. Chignolin, the smallest protein known today, is used as target compound.

# “ Computation of atmospheric delay corrections on GPUs”

Thomas Hobiger (NICT) Ryuichi Ichikawa (NICT), Yasuhiro Koyama (NICT),
Tetsuro Kondo (NICT; Ajou Univ, Korea)

Troposphere delays are one of the major error sources of space geodetic and remote sensing techniques. Unlike ionospheric delays, which have a dispersive characteristics and thus can be canceled out by dual-frequency measurements, atmospheric propagation effects have to be modeled within the post-processing or model adjustment stage. Atmosphere delays within remote sensing techniques are neglected or considered only by simple models, which are not capable to represent larger areas or complex weather phenomena. Since numerical weather models have undergone tremendous improvements concerning accuracy and spatial resolution it has become feasible to utilize this information to compute ray-traced troposphere delays which can be applied as corrections for space geodetic techniques. In order to reduce computation time, simplified algorithms have been developed, which yield similar results, differing by less than a millimeter for usual observing geometries. Nevertheless, even these fast algorithms are not capable to meet the requirements of real-time applications, nor can they compute millions of rays, as needed for remote sensing applications, in reasonable time. Since ray-tracing can be done in parallel, the obvious step of carrying out the computation of troposphere delays on a GPU has been investigated. Given that modern graphic chips support full double precision representation one could expect nearly identical results as obtained on the CPU. After simulating a random observation geometry for a ground GPS receiver, the performance of the ray-tracing algorithms on the GPU respectively CPU can be compared with each other. As long as the number of rays is small, ray-tracing on the CPU is faster, but already for about 1000 rays the GPU starts to perform much better than the fastest CPU algorithms. If the number of rays exceeds 10,000, the GPU solution is at least 15 times faster than the single-core CPU computations. All space geodetic and remote sensing techniques, which are utilizing signals within the microwave frequency band encounter atmosphere delays, which have to modeled within post-processing. Thus, ray-tracing on the GPU seems to be the way to go in the future, providing both, accurate troposphere slant delays and high computational performance, which delivers corrections to the user in near real-time and reduces processing time for a large number of rays.

References:

Hobiger T.,Ichikawa R., Koyama Y., Kondo T., Fast and accurate ray-tracing algorithms

for real-time space geodetic applications using numerical weather models, J. Geophys. Res., 113, D20302, doi: 10.1029 /2008JD010503, 2008.
Hobiger T.,Ichikawa R., Koyama Y., Kondo T., Computation of troposphere slant delays on a GPU, IEEE Trans. Geosci. Remote Sens., 10.1109/TGRS.2009.2022168, in print, 2009.

# “GUI Code Developing Environment for GPGPU on RICC : RIKEN Integrated Cluster of Clusters”

Shigeho, Noda(RIKEN) ,Hideaki Komatsu(IBM), Masana Murase(IBM)

## Aug. 29th(Sat) 9:30-13:10
## Oral

“GPU Based Acceleration of First Principles Calculation”

Hidekazu Tomono (Meiji Univ.) Masaru Aoki(Shizuoka Sngyo Univ),
Toshiaki Iitaka(RIKEN), Kazuo Tsumuraya (Meiji Univ.)

We produced an original GPGPU interface to first principles electronic Structure calculations. It makes FFT calculation, which is the most time-consuming part, 10 times faster. As the result, the total computation time is reduced to 15% of that of the CPU.

# “A real-time GPS software receiver running on the GPU “

Thomas, Hobiger (NICT), Tadahiro Gotoh(NICT),
Jun Amagai (NICT), Tetsuro Kondo (NICT; Ajou Univ., Korea),
Yasuhiro Koyama (NICT)

Global positioning service (GPS) software radios have become an attractive and cheap alternative to hardware receivers. Such software implementations can be easily adopted to handle new GPS signals and allow for rapid testing of innovative algorithms. Although modern CPUs are more and more evolving to multi-core processors, the performance from a single chip is usually not sufficient to fulfill the needs of real-time GNSS tracking. On the other hand, Graphic Processing Units (GPUs) have undergone a tremendous development in the recent years and deployed electronic components have become powerful parallel systems which can be utilized for a variety of applications. Since many of the signal processing stages of a GPS receiver can be parallelized the GPU seems to be a sophisticated candidate for the implementation of software defined radio. Hardware components are only required for down-conversion and digitization of the RF signals. Our system allows to read the A/ D converted data-stream via USB 2.0 feeding it to ringbuffers on the CPU memory before data are copied to the GPU for processing. Our development differs from hardware-receiver architectures in various ways. First, we are trying to deal with the complete correlation function rather than computing only three different time lags (early/prompt/late” scheme) as done on hardware-receivers. This has several advantages for tracking of weaker signals as well as it allows a reduction of multi-path effects. Secondly, our architecture maintains the cross-spectrum throughout all processing stages, which enables us to apply exact filters, rather than using FIR implementations as implemented in hardware receivers. Moreover, our development enables us to run the receiver in closedloop as well as in open-loop mode. Whereas the first mode is usually selected for positioning applications, the open-loop mode can be chosen to track weak signals, a situation which occurs e.g. when GPS rec eivers are utilized for atmosphere sounding. The correlation engines, which are the core of each GPS receivers are implemented with CUFFT allowing us to handle data with sampling rates of up to 32 Msps in real-time. Since the parallel performance of the libraries decreases with the FFT lengths, real-time operation for higher sampling rate is limited to fewres atellites.

Therefore, our presentation will deal with all time-critical stages of such a GPS software receiver. Moreover we are going to present results from real-time GPS tracking and we

will compare these outcomes with those obtained from a commercial hardware receiver which shares the antenna with the software receiver.

References

Hobiger T., T. Gotoh, J. Amagai, Y. Koyama and T. Kondo, A GPU based real-time GPS software receiver, GPS Solutions, doi:10.1007/s10291-009-0135-2, in print, 2009.

“Multi-precision linear equation solver for lattice QCD”

Michael Clark (Harvard ,Boston)

The exponential growth of floating point power in graphics processing units (GPUs), together with their low cost, has given rise to an attractive platform upon which to deploy lattice QCD calculations. The large ratio of raw floating point operations per second to memory bandwidth that is characteristic of GPUs and the relative cost of using double precision arithmetic over that of single necessitates unique algorithmic design choices. In this talk I shall describe these choices, and compare the merits of GPUs against more traditional HPC platforms.

# “A mixed-precision matrix multiplication library for GPUs and its application to Quantum Chemistry calculations”

Mark Watson (Harvard)

Quantum Chemistry is typically concerned with the first principles solution of the Schrodinger equation for atoms and molecules -- a task requiring vast amounts of computational resources if high accuracy is desired. In this talk, I will discuss a new library (MGEMM) for accelerating matrix-matrix multiplications using graphical processing units (GPUs) and its application to second-order Moller-Plesset perturbation theory (MP2) calculations. The new library is designed within a heterogeneous computing model, whereby fast single-precision matrix multiplications involving `small' matrix elements are performed on the GPU (using the optimized CUBLAS library) and augmented with double-precision operations involving `large' matrix elements on the CPU. In this manner, it is possible to control the maximum errors arising from the use of MGEMM compared to full double-precision results and to balance them with the considerable speedups obtained from leveraging the GPUs. We demonstrate the accuracy and efficiency of the new library using model examples and MP2 correlation energy calculations of a variety of molecules.

”GPU acceleration of ab-initio quantum Monte Carlo simulations and its application to molecular crystals”

Kenta Hongo (Harvard)

Quantum Monte Carlo (QMC) methods are some of the most accurate and efficient methods in state-of-the-art ab-initio calculations for studying the electronic structure of atoms, molecules, and solids. This is mainly because electron correlation effects can be taken into account using an explicitly correlated many electron wave function and the computational cost scales moderately as $N^3$ - $N^4$, where N is the number of electrons in the system. Moreover, the intrinsic parallelism of QMC algorithm, make it an ideal candidate for acceleration in the many-core pradigm. This enables one to leverage the advantage of Graphical Processing Units (GPUs). In QMC algorithm, the most time consuming part is the evaluation of many-electron wave function, which is to be executed on the GPUs. In this study, we apply a GPU-enabled QMC code to simple crystals such as diamond and boron-nitride, observing speedups of approximately 10x relative to a single CPU core. We also study polymorphism of molecular crystals, which can be calculated with the help of GPUs. All the present calculations are performed using the program suit QMCPACK.

“Computational Aspects of ab-initio QMC calculation”

Ryo Maezono (JAIST), Tomoharu Terashima(JAIST)

Large scale Monte Carlo simulation technique has been developed for decades in quantum many-body theories of electrons such as ab-initio condensed matter physics and quantum chemistry [1,2]. Recent high performance computing and developments in numerical optimization implementations have brought about growing feasibility of such methods. In the talk topics about the computational/technical aspects of ab-initio QMC calculations are presented. The contents would include a brief introduction of the research field, the description of computational structure, MPI parallel implementation of the code, and the construction of our hand-made PC cluster with several hundreds cores. Recent researches in our group are also presented, such as the evaluation of new random number generators working on ab-initio QMC calculations [3], the modification of the QMC code for GPU computation.

[1] "Quantum Monte Carlo simulations of solids",W.M.C. Foulkes, L. Mitas, R.J. Needs, and G. Rajagopal, Rev. Mod.Phys. 73, 33 (2001).

[2] "Equation of state and Raman frequency of diamond from quantum Monte Carlo simulations", Ryo Maezono et.al., Phys. Rev. Lett., 98, 025701:1-4 (2007).

[3] "Random number generators tested on quantum Monte Carlo simulations" Ryo Maezono, Kenta et.al. and Ken-ichi Miura, to be published.